

PRIORITIES GENERATION AND MANAGEMENT

FIELD OF THE INVENTION

This invention relates generally to the prioritization of documents by measures of
5 importance, and more particularly to the generation and management of such priorities.

BACKGROUND OF THE INVENTION

Electronic mail programs have become a popular application among computer
users. Especially with the advent of the Internet, exchanging email has almost become a
10 reason why people purchase computers for personal reasons, and within many corporate
environments, email has become the standard manner by which coworkers exchange
information. However, with the increasing popularity of email, shortcomings have
become apparent.

Chief among these shortcomings is that many users now face a deluge of email
15 every day, such that the capability of being able to send and receive email has almost
become a hindrance to their day-to-day ability to get their job done, as opposed to being
an asset. Some users report receiving over 100 email messages a day. With such large
numbers of email, it is difficult to manage the email, such that the users read the most
important messages first.

20 Limited solutions to this problem have been attempted in the prior art. Prior art
exists for attempting to curtail the amount of junk email that users receive. Moreover,
some electronic mail programs allow for the generation of rules that govern how an email

is managed within the program. For example, all emails from certain coworkers are placed in a special folder.

These limited solutions, however, do not strike at the basic problem behind email. With so much email being received, it would be most useful for a user to be able to have his or her computer automatically prioritize the email by importance or review urgency, and perform actions based on that prioritization. For these and other reasons, there is a need for the present invention.

SUMMARY OF THE INVENTION

10 The invention relates to generating and managing priorities for texts such as email messages. A priority for a text is generated based on a classifier, such as by determining the likelihood that the text is of high priority. The classifier may be a Bayesian classifier, a support vector machine, or another type of classifier. The classifier may consider features such as the structural relationship between the user and the sender of the text, as
15 well as the time of any event referenced in the text. The expected loss of non-review of the text at a current time is determined based on the priority, as well as the expected cost of outputting the text at the current time. The user is alerted on a mobile device in response to determining that the expected loss is greater than the expected cost, in accordance with a profile.

20 The current profile is selected from a number of profiles, each of which is editable by the user to reflect a different context. The profiles are schedulable on a per-day and by-time basis, and can be locked to remain active unless subsequently unlocked. Each profile has a priority threshold settable by the user through a coarsening mechanism. A

chunk setting of each profile controls delivery of the text. For example, the text may be delivered to the mobile device only in conjunction with one or more other texts.

Alternatively, the text may be delivered to the mobile device only when a specified period has expired. Furthermore, the text may be delivered to the mobile device only if a

- 5 primary device of the user, such as a desktop computer, is idle for a specified period of time.

The text may be formatted prior to delivery to the mobile device. Formatting can include compression and fragmentation. In the case of the former, the text is compressed according to a specified compression setting that is settable by the user through a

- 10 coarsening mechanism. In the case of the latter, there is also a user-settable fragmentation setting. Where the text is email, the sender of the text may be sent an indication that the text has been delivered to the mobile device of the user. Furthermore, the sender may be sent an indication that the user is away from his or her primary device. Alerting the user may also consider the user's calendar. Each event on the calendar has
15 an associated tag that specifies the degree to which the user is amenable to being interrupted during the event. Other aspects and embodiments of the invention, beyond those described here, will become apparent by reading the following detailed description, and by reference to the drawings.

20

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a diagram of an operating environment in conjunction with which embodiments of the invention may be practiced.

FIG.2 is a diagram showing explicit and implicit training of a text classifier.

FIG. 3 is a diagram showing how a priority for a text is generated by input to a text classifier.

FIG. 4(a) is a diagram of a scheme according to which the priority of a text can be classified.

5 FIG. 4(b) is a diagram of another scheme according to which the priority of a text can be classified.

FIG. 5(a) is a graph showing linear cost functions of high, medium, and low priority texts.

FIG. 5(b) is a graph showing a non-linear cost function for a text.

10 FIG. 6 is a diagram showing classes of evidence that can be used to make an inference about a user's activity.

FIG. 7 is a diagram showing a Bayesian network that can be used for inferring a user's activity.

15 FIGs. 8-10 are influence diagrams showing how decision models can be utilized to make the decision as to how and when to alert a user to a message.

FIG. 11 is a flowchart of a method.

FIG. 12 is a diagram of a system.

FIGs. 13(a)-13(c) are diagram of systems other than that of FIG. 12.

20 FIG. 14(a) is a diagram of a user interface via which alert criteria can be modified.

FIGs. 14(b) and 14(c) are diagrams of a user interface via which routing criteria can be modified.

FIG. 14(d) is a diagram of a user interface via which predetermined criteria for alerting can be modified.

FIGs. 15-20 are diagrams of a user interface for routing priorities to a mobile device different from that described in conjunction with FIG. 14(a).

5 FIG. 21 is a diagram of a user interface in which the user can specify that the sender of a message to the user be informed when the user is expected to return.

FIG. 22 is a diagram showing the classifier-training process, and real-time classification.

10

DETAILED DESCRIPTION OF THE INVENTION

15

In the following detailed description of exemplary embodiments of the invention, reference is made to the accompanying drawings, which form a part hereof, and in which is shown by way of illustration specific exemplary embodiments in which the invention may be practiced. These embodiments are described in sufficient detail to enable those skilled in the art to practice the invention, and it is to be understood that other embodiments may be utilized and that logical, mechanical, electrical, and other changes may be made without departing from the spirit or scope of the present invention. The following detailed description is, therefore, not to be taken in a limiting sense, and the scope of the present invention is defined only by the appended claims.

20

Operating Environment

Referring to FIG. 1, a diagram of the hardware and operating environment in conjunction with which embodiments of the invention may be practiced is shown. The description of FIG. 1 is intended to provide a brief, general description of suitable computer hardware and a suitable computing environment in conjunction with which the invention may be implemented. Although not required, the invention is described in the general context of computer-executable instructions, such as program modules, being executed by a computer, such as a personal computer. Generally, program modules include routines, programs, objects, components, and data structures that perform particular tasks or implement particular abstract data types.

Moreover, the invention may be practiced with other computer system configurations, including hand-held devices, multiprocessor systems, microprocessor-based or programmable consumer electronics, network PC's, minicomputers, mainframe computers, and the like. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

The exemplary hardware and operating environment of FIG. 1 for implementing the invention includes a general purpose computing device in the form of a computer 20, including a processing unit 21, a system memory 22, and a system bus 23 that operatively couples various system components include the system memory to the processing unit 21. There may be only one or there may be more than one processing unit 21, such that the processor of computer 20 comprises a single central-processing unit (CPU), or a plurality

of processing units, commonly referred to as a parallel processing environment. The computer 20 may be a conventional computer, a distributed computer, or any other type of computer.

- The system bus 23 may be any of several types of bus structures including a
- 5 memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. The system memory may also be referred to as simply the memory, and includes read only memory (ROM) 24 and random access memory (RAM)
25. A basic input/output system (BIOS) 26, containing the basic routines that help to transfer information between elements within the computer 20, such as during start-up, is
- 10 stored in ROM 24. The computer 20 further includes a hard disk drive 27 for reading from and writing to a hard disk, a magnetic disk drive 28 for reading from or writing to a removable magnetic disk 29, and an optical disk drive 30 for reading from or writing to a removable optical disk 31.

- The hard disk drive 27, magnetic disk drive 28, and optical disk drive 30 are
- 15 connected to the system bus 23 by a hard disk drive interface 32, a magnetic disk drive interface 33, and an optical disk drive interface 34, respectively. The drives and their associated computer-readable media provide nonvolatile storage of computer-readable instructions, data structures, program modules, and other data for the computer 20. Any type of computer-readable media which can store data that is accessible by a computer,
- 20 such as magnetic cassettes, flash memory cards, digital video disks, Bernoulli cartridges, random access memories (RAMs), read only memories (ROMs), and the like, may be used in the exemplary operating environment.

A number of program modules may be stored on the hard disk, magnetic disk 29, optical disk 31, ROM 24, or RAM 25, including an operating system 35, one or more application programs 36, other program modules 37, and program data 38. A user may enter commands and information into the personal computer 20 through input devices

5 such as a keyboard 40 and pointing device 42. Other input devices (not shown) may include a microphone, joystick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 21 through a serial port interface 46 that is coupled to the system bus, but may be connected by other interfaces, such as a parallel port, game port, or a universal serial bus (USB). A monitor 47 or other

10 type of display device is also connected to the system bus 23 via an interface, such as a video adapter 48. In addition to the monitor, computers typically include other peripheral output devices, such as speakers and printers.

The computer 20 may operate in a networked environment using logical connections to one or more remote computers, such as remote computer 49. These

15 logical connections are achieved by a communication device coupled to or a part of the computer 20. The remote computer 49 may be another computer, a server, a router, a network PC, a client, a peer device or other common network node, and typically includes many or all of the elements described above relative to the computer 20, although only a memory storage device 50 has been illustrated in FIG. 1. The logical

20 connections depicted in FIG. 1 include a local-area network (LAN) 51 and a wide-area network (WAN) 52. Such networking environments are commonplace in office networks, enterprise-wide computer networks, intranets and the Internet, which are all types of networks.

When used in a LAN-networking environment, the computer 20 is connected to the local network 51 through a network interface or adapter 53, which is one type of communications device. When used in a WAN-networking environment, the computer 20 typically includes a modem 54, a type of communications device, or any other type of 5 communications device for establishing communications over the wide area network 52, such as the Internal. The modem 54, which may be internal or external, is connected to the system bus 23 via the serial port interface 46. In a networked environment, program modules depicted relative to the personal computer 20, or portions thereof, may be stored in the remote memory storage device. It is appreciated that the network connections 10 shown are exemplary and other means of and communications devices for establishing a communications link between the computers may be used.

Generating Measures for Prioritization of Text Documents: Classification

In this section of the detailed description, the generation of a priority for text 15 documents such as an email is described. The generation of priorities for texts as described can then be used in methods, systems, and computer-readable media. The description in this section is provided in conjunction with FIG. 2 and FIG. 3, the former which is a diagram showing explicit and implicit training of a text classifier, and the latter which is a diagram showing how a priority for a text is generated by input to a text 20 classifier. The description is also provided in conjunction with FIGs. 4(a) and 4(b), which are diagrams of different schema according to which the priority of a text can be classified, and in conjunction with FIGs. 5(a) and 5(b), which are graphs showing different cost functions that may be applicable depending on text type.

Referring first to FIG. 2, the text classifier 200 can be trained both explicitly, as represented by the arrow 202, and implicitly, as represent by the arrow 204. The explicit training represented by the arrow 202 is usually conducted at the initial phases of constructing the text classifier 200, while the implicit training represented by the arrow 5 204 is usually conducted after the text classifier 200 has been constructed, to fine tune the classifier 200.

The text classifier 200 can be a Bayesian classifier, a naïve-Bayesian classifier, a support vector machine (SVM) classifier, or another type of classifier. Text classification methodology based on a Bayesian learning approach is specifically described in the 10 reference M. Sahami, S. Dumais, D. Heckerman, E. Horvitz, A Bayesian Approach to Junk E-Mail Filtering, AAAI Workshop on Text Classification, July 1998, Madison, Wisconsin, AAAI Technical Report WS-98-05. Text classification methodology based on an SVM approach is specifically described in the following references:

- The coassigned patent, U.S. Pat. No. 5,864,848, issued January 26, 1999;
- 15 • The previously filed and coassigned case entitled "Methods and Apparatus for Building a Support Vector Machine Classifier," serial no. 09/055,477, filed on April 6, 1998; and,
- The reference J. Platt, Fast Training of Support Vector Machines using Sequential Minimal Optimization, MIT Press, Baltimore, MD, 1998.

20 Specific description is made herein with reference to an SVM classifier. Other text classification approaches include Bayesian networks, decision trees, and probabilistic classification models assuming different patterns of independence. Text classification as

used herein also is inclusive of statistical regression that is used to develop models of priority.

As shown in FIG. 2, the explicit training of the text classifier 200 as represented by the arrow 202 includes constructing the classifier in 206, including utilizing feature selection. In general, Support Vector Machines build classifiers by identifying a hyperplane that separates a set of positive and negative examples with a maximum margin. In the linear form of SVM, the margin is defined by the distance of the hyperplane to the nearest positive and negative cases for each class. Maximizing the margin can be expressed as an optimization problem. A post-processing procedure described in the Platt reference is used that employs regularized maximum likelihood fitting to produce estimations of posterior probabilities. The method fits a sigmoid to the score that is output by the SVM classifier.

In the explicit training, the text classifier is presented with both time-critical and non-time-critical texts, so that it may be able to discriminate between the two. This training set may be provided by the user, or a standard training set may be used. Given a training corpus, the text classifier first applies feature-selection procedures that attempt to find the most discriminatory features. This process employs a mutual-information analysis. Feature selection can operate on single words or higher-level distinctions made available, such as phrases and parts of speech tagged with natural language processing. That is, the text classifier 200 can be seeded with specially tagged text to discriminate features of a text that are considered important.

Feature selection for text classification typically performs a search over single words. Beyond the reliance on single words, domain-specific phrases and high-level

patterns of features are also made available. Special tokens can also enhance classification. The quality of the learned classifiers for email criticality can be enhanced by inputting to the feature selection procedures handcrafted features that are identified as being useful for distinguishing among email of different time criticality. Thus, during 5 feature selection, single words as well as special phrases and symbols that are useful for discriminating among messages of different levels of time criticality are considered.

Tokens and patterns of value in identifying the criticality of messages include such distinctions as, and including Boolean combinations thereof:

Information in the Header

- 10 **To: field (Recipient information)**
Addressed just to user
Addressed to only a few people including user
Addressed to an alias with a small number of people
15 Addressed to several aliases with a small number of people
Cc:’d to user
Bcc:’d to user
- 20 **From: field (Sender information)**
Names on pre-determined list of important people, potentially segmented into a variety of classes of individuals, e.g., Family members, Friends
Senders identified as internal to the user’s company organization
Information about the structure of organizational relationships relative to the user drawn from an online organization chart
25 Managers user reports to
Managers of the managers of users
People who report to the user
External business people
- 30 **Past tense**
These include descriptions about events that have already occurred such as:

We met
meeting went
35 happened
got together
took care of
meeting yesterday
- 40 **Future tense**
Tomorrow
This week
Are you going to
When can we

Looking forward to
Will this
Will be

5 **Meeting and coordination**

Get together
Can you meet
Will get together
Coordinate with
10 Need to get together
See you
Arrange a meeting
Like to invite
Be around
15

Resolved dates

Future vs. past dates and times indicated from patterns of text used to state dates and times explicitly or typical abbreviations

20 On 5/2

Questions

Words, phrases adjacent to question marks ?

25

Indications of personal requests:

Can you
Are you
Will you
30 you please
Can you do
Favor to ask
From you

35

Indications of need:

I need
He needs
She needs
I'd like
40 It would be great
I want
He wants
She wants
Take care of

45

Time criticality

happening soon
right away
deadline will be
50 deadline is
as soon as possible
needs this soon
to be done soon
done right away
55 this soon

by [date]
by [time]

- 5 **Importance**
is important
is critical
Word, phrase + !
Explicit priority flag status (low, none, high)
- 10 **Length of message**
Number of bytes in component of new message
- 15 **Signs of Commercial and Adult-Content Junk email**
Free!!
Word + !!!
Under 18
Adult's only
Percent of capitalized words
Percent nonalphanumeric characters
etc.
- 20

Furthermore, still referring to FIG. 2, implicit training of the text classifier 200, as represented by the arrow 204, can be conducted by continually watching the user work in 210. The assumption is that as users work, and lists of mail are reviewed, time-critical 25 messages are read first, and low-priority messages are reviewed later, or just deleted. That is, when presented with a new email, the user is watched to determine whether or she immediately opens the email, and in what order, deletes the email without opening, and/or replies to the email right away. Thus, the text classifier is such that a user is continually watched while working, and the classifier is periodically refined by training 30 in the background and updated for enhancing the real-time decision-making. Background methods for building classifiers can extend from those that update the classifier with every new training message.

Alternatively, larger quantities of messages are gathered and new filters are created in a batch process, either per a daily schedule, per the number of new quantities 35 of messages admitted to the training set, or combinations. For each message inputted into

the classifier, a new case for the classifier is created. The cases are stored as negative and positive examples of texts that are either high or low priority. What can be recognized are low, medium, and high urgency classes and such that the probabilities of membership in each of these classes are used to build an expected criticality. Larger numbers of 5 criticality classes can be used to seek higher resolution. For example, as shown in FIG. 22, a training set of messages 5500 can be initially used to train the classifier, such that real-time classification is accomplished, as indicated by 5502, based on the initial training.

In addition, a simple estimation of a number can be done directly by watching a 10 user interact with email, rather than labeling the case as one of a specific small set of folders. A classifier can be continued to be updated but have a moving window, where only cases of messages or documents that are newer than some age are considered, specified by the user.

The constant rate of loss associated with the delayed review of messages is 15 referred to as the expected criticality (EC) of a message,

$$EC = \sum_i C^d(H_i) p(H_i | E^d)$$

where C is a cost function, E is an event, H is the criticality class of the email, and the EC is the sum over the likelihood of the classes weighted by the rate of lost described by the cost function C for each of the potential classes.

20 Referring next to FIG. 3, a text, such as an email message, 300 is input into the text classifier 200, which based thereon generates a priority 302 for the text 300. That is, the text classifier 200 generates a priority 302, measured as a percentage from 0 to 100%.

This percentage is a measure of the likelihood that the text 300 is of high priority, based on the previous training of the classifier 200.

It is noted that as has been described, the text classifier and the priority are based on a scheme where each email in the training phase is construed as either high priority or 5 low priority. The priority generated by the text classifier is a measure of the likelihood of the text being analyzed is of high priority. This scheme is specifically shown by reference to FIG. 4(a). The text classifier 200 is trained by a group of texts 400 that are high priority and a group of texts 402 that are low priority. A text to be analyzed 400 is input into the classifier 200, which outputs a scalar number 406 measuring the likelihood 10 that the text being analyzed is of high priority.

For example, referring to FIG. 4(b), a diagram showing a scheme where texts are divided into low, medium, and high priority is shown. The text classifier 200 is trained by a group of texts 400 that are high priority and a group of texts 402 that are low priority, and by a group of texts 450 that are medium priority. Thus, a text to be analyzed 15 400 is input into the classifier 200, which outputs a scalar number 406, that can measure the likelihood that the text being analyzed is of high priority, if so desired, or medium priority or low priority. The classifier 200 is also able to output a class 452, which indicates the class of low, medium, or high priority that the text 404 most likely falls into. Further classes can also be added if desired.

20 The invention is not limited to the definition of priority as this term is used by the text classifier to assign such priority to a text such as an email message. Priority can be defined in terms of a loss function. More specifically, priority is defined in terms of the expected cost in lost opportunities per time delayed in reviewing the text after it has be

received. That is, the expected lost or cost that will result for delayed processing of the text. This loss function can further vary according to the type of text received.

For example, the general case is shown in FIG. 5(a), which is a graph of linear cost functions dependent on the priority of a text. In the graph 500, as time increases, the 5 cost of not having reviewed a text also increases. However, the cost increases more for a high priority message, as indicated by the line 502, as compared to a medium priority message, as indicated by the line 504, or a low priority message, as indicated by the line 506. That is, the high priority line 502 may have a slope of 100, the medium priority line 504 may have a slope of 10, and the low priority line 502 may have a slope of one. These 10 slope values can then be used by the text classifier to assist in assigning a priority to a given text, for example, by regression analysis.

Some messages, however, do not have their priorities well approximated by the use of a linear cost function. For example, a message relating to a meeting will have its cost function increase as the time of the meeting nears, and thereafter, the cost function 15 rapidly decreases. That is, after the meeting is missed, there is not much generally a user can do about it. This situation is better approximated by a non-linear cost function, as shown in FIG. 5(b). In the graph 550, the cost function 554 rapidly increases until it reaches the time of the meeting demarcated by the line 552, after which it rapidly decreases. Depending on a message's type, the cost function can be approximated by one 20 of many different representative cost functions, both linear and non-linear.

Thus, as has been described, the priority of a text can be just the likelihood that it is of high priority based on output of a text classifier, or the most likely priority class it

falls into, also based on the output of the text classifier. Alternatively, an expected time criticality of each text, such as an email message, is determined. This can be written as

$$EL = \sum_i^n p(critical_i)C(critical_i)$$

where EL is the expected loss, $p(critical_i)$ is the probability that a text has the criticality i ,

- 5 $C(critical_i)$ is the cost function for text having the criticality i , and n is the total number of criticality classes minus one. The cost functions may be linear or non-linear, as has been described. In the case where the function is linear, the cost function defines a constant rate of loss with time. For non-linear functions, the rate of loss changes with delayed review or processing of the text and can increase or decrease, depending on the amount of
- 10 delay.

In the case where $n=1$, specifying that there are only two priority classes low and high, the expected loss can be reformulated as

$$EC = p(critical_{high})C(critical_{high}) + [1 - p(critical_{low})]C(critical_{low})$$

where EC is the expected criticality of a text. Furthermore, if the cost function of low

- 15 criticality messages is set to zero, this becomes

$$EC = p(critical_{high})C(critical_{high})$$

The total loss until the time of review of a text can be expressed as the integration of the expressed criticality, or,

$$EL = \int_0^t p(critical_{high})C(critical_{high}, t)dt$$

- 20 where t is the time delay before reviewing the document.

Other measures that accord a valuable metric for ranking documents, such as email messages, by importance. While the discussion above focused on priority as time

criticality, other notions of “importance” can also be trained. For example, this can be accomplished by labeling a set of training folders: “High Importance” all the way down to “Low Importance” a measure of “expected importance” can be computed. Another metric can be based on the use of the semantic label, “messages that I would wish to hear about within 1 day while traveling” and to compute a measure for prioritizing messages for forwarding to traveling user.

5

Furthermore, one utilized metric is urgency or time-criticality, as it has clear semantics for decision making, triage, and routing. In this case, the classes are labeled according to different levels of urgency and compute an expected urgency for each

10 message from the probabilities inferred that the message is in each class.

Extensions to Criticality Classification

Extensions to criticality classification, as described in the previous section, can also be accommodated. For instance, classification can include an automatic search for combinations of high-payoff features within or between classes of features. As an example, combinations of special distinctions, structures, and so on, with words that have been found to be particularly useful for certain users can be searched for and used in the classification process. A combination of two features is referred as a doublet, whereas a combination of three features is referred to as a triplet, and so on. The combination of features can allow for better classification to result. Classification can also be improved with the use of incremental indexing that uses a moving window in the classifier. This allows the classifier to be continually refreshed, as old data is timed out, and new data is brought in.

15

20

Classification can also be based on the determination of the date and time of the event specified in a message. This determination is used to assign features to the message that can be utilized by the classifier. For example, the features assigned may include: today within four hours, today within eight hours, tomorrow, this week, this month, and next month and beyond. This enables the classifier to have improved accuracy with respect to the messages that it is classifying. In general, classification can be based on the time of the referenced event, considering whether the event is in the future or has past. With respect to future events, classification thus considers the sender's reference to a time in the future when the event is to occur.

- 10 Other new features can also be integrated into the classification process. For example, an organization chart can be utilized to determine how important a message is by the sender's location within the chart. Linguistic features may be integrated into the classifier. To accommodate different languages, the features may be modified depending on the origin of the sender, or the language in which the message is written.
- 15 Classification may be different depending on different folders in which messages are stored, as well as other scaling and control rules. Besides email, classification can be performed on instant messages, and other sources of information, such as stock tickers, and so on.

In general, the sender-recipient structural relationship may be considered in the classification process. If the user is the only recipient of a message, then this message may be considered as more important than a message sent to a small number of people. In turn, a message sent to a small number of people may be more important than a message on which the user is only blind-copied (bcc'ed) or carbon-copied (cc'ed). With

respect to the sender, criticality may be assigned based on whether the sender's name is recognized. Criticality may also be assigned depending on whether the sender is internal or external to the organization of which the user is a part.

- Other distinctions that may be considered in classification include the length of
- 5 the message, whether questions have been detected, and whether the user's name is in the message. Language associated with time criticality may increase the message's importance. For example, phrases such as "happening soon," "right away," "as soon as possible," "ASAP," and "deadline is," may render the message more critical. Usage of past tense as compared to future tense may be considered, as well as coordinative tasks
- 10 specified by phrases such as "get together," "can we meet," and so on. Evidence of junk mail may lower the priority of a message. Predicates representing key combinations, such as a short question from a sender proximate to the user in the organization chart, may also be considered in the classification process.

15 Determining When to Alert the User

- In this section of the detailed description, described is provided as to determining when to alert the user of a high-priority text, for example, a text that has a likelihood of being high priority greater than a user-set threshold, or greater than a threshold determined by decision-theoretic reasoning. That is, beyond knowing about time-critical
- 20 messages, it is also important to decide when to alert a user to time-critical messages if the user is not directly viewing incoming email. In the general case, a cost of distracting the user from the current task being addressed to learn about the time-critical message is determined.

Alternatively, different policies for alerting and notification can be used. Two policies include:

- Put a user-specified upper bound on the total loss. This policy would specify that the system should generate an alert when the total loss associated with the delayed review of a message exceeds some pre-specified “tolerable” loss x .
- Another policy is a cost-benefit analysis based on more complete decision-theoretic analysis, such as $NEVA = EVTA - ECA - TC$, where $NEVA$ is the net expected value of alerting, $EVTA$ is the expected value of alerting, ECA is the expected cost of alerting, and TC is the transmission cost.

In general, a user should be alerted when a cost-benefit analysis suggests that the expected loss the user would incur in not reviewing the message at time t is greater than the expected cost of alerting the user. That is, alerting should be conducted if

$$EL - EC > 0$$

where EL is the expected loss of non-review of the text at a current time t , and EC is the expected cost of alerting the user of the text at the current time t . The expected loss is as described in the previous section of the detailed description.

However, this formulation is not entirely accurate, because the user will often review the message on his or her own in the future. Therefore, in actuality, the user should be alerted when the expected value of alerting, referred to as $EVTA$, is positive. The expected value of alerting should thus consider the value of alerting the user of the text now, as opposed to the value of the user reviewing the message later on his or her own, without alert, minus the cost of alerting. This can be stated as

$$EVA = EL_{\text{alert}} - EL_{\text{no-alert}} - EC$$

where EL_{alert} is the expected loss of the user reviewing the message if he or she were to review the message now, upon being alerted, as opposed to $EL_{\text{no-alert}}$, which is the expected loss of the user reviewing the message on his or her own at some point, without being alerted, minus EC , the expected cost of alerting based on a consideration of distraction and on the direct cost of the transmitting the information.

Furthermore, information from several messages is grouped together into a single compound alert. Reviewing information about multiple messages in an alert can be more costly than an alert relaying information about a single message. Such increases in distraction can be represented by making the cost of an alert a function of its informational complexity. It is assumed that the EVA of an email message is independent of the EVA of the other email messages. $EVA(M_i, t)$ is used to refer to the value of alerting a user about a single message M_i at time t and $ECA(n)$ is used to refer to the expected cost of relaying the content of n messages. Thus, multiple messages can be considered by summing together the expected value of relaying information about a set of n messages,

$$NEVA = \sum_{i=1} EVA(M_i, t) - ECA(n).$$

It is also noted that in order to determine the expect cost of alerting, it is useful to infer or directly access information about whether the user is present or is not present.

Sensors can be used that indicate when a user is in the office, such as infrared sensors and pressure sensors. However, if such devices are not available, a probability that a user is in the office can be assigned as a function of user activity on the computer, such as the time since last observed mouse or keyboard activity. Furthermore, scheduling

information available in a calendar can also be made use of to make inferences about the distance and disposition of a user, to consider the costs of forwarding messages to the user by different means.

- It is also important to know how busy the user is in making decisions about
- 5 interrupting the user with information about messages with high time criticality. It is reasoned about whether and the rate at which a user is working on a computer, or whether the user is on the telephone, speaking with someone, or at a meeting at another location.
- Several classes of evidence can be used to assess a user's activity or his or her focus of attention, as shown in FIG. 6. A Bayesian network can then be used for performing an
- 10 inference about a user's activity. An example of such a network is shown in FIG. 7.

In general, a decision should be made as to when and how to alert users to messages and to provide services based on the inference of expected criticality and user activity. This decision is made by utilizing decision models. FIGs. 8-10 are influence diagrams showing how such decision models can be utilized to make this decision.

15 Specifically, FIG. 8 displays a decision model for decisions about interrupting a user, considering current activity, expected time criticality of messages, and cost of alerting depending on the modality. FIG. 9 also includes variables representing the current location and the influence of that variable on activity and cost of the alternate messaging techniques. Finally, FIG. 10 is further expanded to consider the costs associated with

20 losses in fidelity when a message with significant graphics content is forwarded to a user without the graphical content being present.

Alternatively, the decision as to when and how to alerts users is made by employment of a set of user-specified thresholds and parameters defining policies on

alerting. User presence can be inferred based on mouse or keyboard activity. Thus, a user can be allowed to input distinct thresholds on alerting for inferred states of activity and nonactivity. Users can input an amount of idle activity following activity where alerting will occur at lower criticalities. If it is determined that the user is not available

- 5 based on the time that no computer activity is seen, then messages are stored, and are reported to the user in order of criticality when the user returns to interact with the computer. Furthermore, users can specify routing and paging options as a function of quantities including expected criticality, maximum expected loss, and value of alerting the user.

- 10 The system may also estimate when the user is expected to return, such that it transmits priorities that are expected to be important before the user is expected to return. This can be accomplished by learning user-present and user-away patterns over time. The user can then set appropriate policies in terms of when he or she is expected to return to the system to review the priorities without being alerted to them. The expected time to
15 return determination by the system may be automatically conveyed to senders of highly urgent messages. In this way, the senders know when the user is expected to return such that he or she can reply to the messages. For example, as shown in the window 5000 of FIG. 21, the user can specify whether to let the sender know when he or she is expected to return as indicated by the circle 5002. The sender may also be informed that his or her
20 message has been conveyed to the user's mobile device, and so on.

Method and System

In this section of the detailed description, a computer-implemented method and a computerized system are described. With respect to the method, the method is desirably realized at least in part as one or more programs running on a computer -- that is, as a program executed from a computer-readable medium such as a memory by a processor of a computer. The program is desirably storable on a machine-readable medium such as a floppy disk or a CD-ROM, for distribution and installation and execution on another computer.

Referring to FIG. 11, a flowchart of a method is shown. In 900, a text to have a priority thereof assigned is received. The text can be an email message, or any other type of text. In 902, the priority of the text is generated, based on a text classifier, as has been described. Thus, 902 includes initially training and continually training the text classifier, as has been described.

The priority of the text is then output in 904. As indicated in FIG. 11, this can include 906, 908, 910, 912, and 914. In 906, an expected loss of non-review of the text at a current time t is determined. This determination considers the expected loss of now-review of the text at a future time, based on the assumption that ultimately the user will review the text him or herself, without being alerted, as has been described. In 908, an expected cost of alerting is determined, as has also been described. If the loss is greater than the cost in 910, then no alert is made at the time t , and the method proceeds back to 906, to redetermine the cost-benefit analysis, at a new current time t . This is done because as time progresses, the expected loss may at some point outweigh the alert cost,

such that the calculus in 910 changes. Upon the expected loss outweighing the alert cost, then an alert to the user is performed in 914, as has been described.

The output of the alert is performed as is now described. A user is alerted on an electronic device based on alert criteria, which indicates when the user should be alerted
5 of a prioritized text. The electronic device on which the user is alerted can be a pager or a cellular telephone.

The method alerts a user on an electronic device, such as a pager or a cellular phone, based on alert criteria that can be made sensitive to information about the location, inferred task, and focus of attention of the user. Such information can be inferred under
10 uncertainty or can be accessed directly from online information sources. The information from an online calendar is made use of to control the criteria used to make decisions about relaying information to a mobile device.

The output of the alert can be performed as is now described. The alert is performed by routing the prioritized text based on routing criteria. Routing of the text
15 can include forwarding the text, or replying to the sender of the text, in the case where the text is email. For example, a sound is played to alert the user to a prioritized document.

Alternatively, an agent or automated assistant is opened. That is, the agent appears on the screen, to notify the user of the prioritized document. Furthermore, the prioritized document is opened, such as being displayed on the screen. The document can receive
20 focus. This can also include sizing the document based on its priority, such that the higher the priority of the document, the larger the window in which it is displayed, and/or centrally locating the document based on its priority.

Referring next to FIG. 12, a diagram of a system is shown. The system includes a program 950 and a text classifier 952. Each of the program 950 and the classifier 952 include a computer program executed by a processor of a computer from a computer-readable medium thereof.

- 5 The program 950 generates a text for input into the text classifier 952. The program includes an electronic mail program that receives email, which then serve as the text. The text classifier 952, based on the text, generates a priority thereof, as has been described. The text classifier 952 can be a Bayesian text classifier, a Support Vector Machine classifier, or another type of classifier. The priority of the text output by the text
10 classifier 952 can then be used in further conjunction with a cost-benefit analysis, as has been described, to effectuate further output and/or alerting based thereon, as has been described.

- Referring next to FIG. 13(a), a diagram of another system is shown. The system of FIG. 13(a) includes an additional component, an alert mechanism 970. Not shown in FIG. 13(a) are the program 950 and the text classifier 952. However, the alert mechanism 970 is operatively and/or communicatively coupled to the latter. The mechanism 970 includes a computer program executed by a processor of a computer from a computer-readable medium thereof.

- As shown in FIG. 13(a), the alert mechanism 970 is communicatively coupled to the Internet 972, which is the network by which the alert mechanism 970 contacts an electronic device to alert the user to a prioritized text, based on an alert criteria. The network is not limited to the Internet 972, however. Thus, the alert mechanism 970 is able to alert the user of a prioritized text via contacting a pager 974, or a cellular phone

976, or other electronic devices capable of receiving information from over a network such as the Internet 972, and which are not shown in FIG. 13(a).

Referring next to FIG. 13(b), a diagram of another system is shown. The system of FIG. 13(b) includes an additional component, a routing mechanism 1970. Not shown 5 in FIG. 13 are the program 950 and the text classifier 952. However, the routing mechanism 1970 is operatively and/or communicatively coupled to the latter. The mechanism 1970 includes a computer program executed by a processor of a computer from a computer-readable medium thereof.

The routing mechanism 1970, as shown in FIG. 13(b), receives a prioritized text, 10 and based on routing criteria, is able to reply to the sender of the text, in which case the mechanism is a replying mechanism. Also based on the routing criteria, the mechanism 1970 is able to forward the text, for example, to a different email address, in which case the mechanism is a forwarding mechanism. The former may be useful when the user wishes to indicate to the sender of a message that the user is not present, and thus may 15 provide the sender with contact information as to how to reach the user. The latter may be useful when the user has email access to a different email address, such as a web-based email address, such that the user wishes to be kept informed of high priority emails at this alternative address.

Referring next to FIG. 13(c), a diagram of another system is shown. The system 20 of FIG. 13(c) includes an additional component, an alerting mechanism 2970. Not shown in FIG. 13(c) are the program 950 and the document classifier 952. However, the alerting mechanism 2970 is operatively and/or communicatively coupled to the latter. The mechanism 2970 includes a computer program executed by a processor of a computer

from a computer-readable medium thereof. The alerting mechanism 2970, as shown in FIG. 13(c), receives a prioritized document, and based on a predetermined criteria, is able to display the document, or play a sound, as has been described. All of the documents that have been received, and that have a priority greater than a predetermined threshold, 5 are displayed as a list.

The system can include other functionality as well. For example, a special priorities-oriented viewer can be provided that acts like a special view onto one's email store, in terms of its ability to filter by priority. The special viewer can allow for summaries of messages to be sorted in a list by priority score. The viewer can also allow 10 a user to sort and view only those messages that remain unread as an option. The special viewer can also allow users to scope the sorting of messages by priority within some scoped time period, and to change the scope or periods being considered. For example, a user can specify that the viewer only display email from today. Alternatively, the user can specify that the list span two days, one week, or all the messages in the in-box. The 15 viewer can also let the user prune from the display messages below a user-specified minimal threshold.

Furthermore, beyond the use of qualitatively different sounds for low, medium, and high priorities, one or more scalar parameters can be used that define the manner by which an alerting sound is rendered. The parameters can be made functions of the 20 inferred priority. Such parameters include variables that dictate something as simple as the volume of the alerting sound, for example, to continuous changes in the modulation or resonance of the sound.

Other functionality includes a simple form can be provided to users to define thresholds among different ranges of uncertainty, and they can specify multiple options involving the automation of the sizing and centering of messages within each range. A “While Away” briefer can be included to give the user a summary of messages that have arrived while a user was away or busy with another application. The system can be instructed to bring up a summary of email directed by priority values when a user returns after being away, or comes back to the special viewer after working with the system in a quiet mode. The automated text summarizer can be controlled to decrease continuously the summarization level of the text of messages as a function of the priority of the document. That is, as documents increase in priority, they are less and less summarized in the summarized view. The priorities can also be used to color or add special annotations, such as priority flags, icons indicating level of priority, and a special priority field itself, to email headers appearing in the display.

Furthermore, a user-defined threshold can be used on the priority assigned to messages to set up a temporary interaction context that is active for some portion of time following an alert or summary that a message has arrived exceeding the threshold. Following an alert, and lasting for the time period that an interaction context is active, predetermined gestures are allowed to give the user access to more details about the message that was associated with the alert. Such gestures include a simple wiggle of the mouse from side to side, for example. As an example, an audio alert may indicate that an incoming message has exceeded some threshold of criticality. The user can then wiggle the mouse quickly from side to side to see details about the message that led to the alert.

The amount of time that such an interaction context is active can be made a function of the priority of the message, or can be user-defined.

Alert Criteria

5 In this section of the detailed description, alert criteria are described. The alert criteria are the criteria that govern when and how a user is alerted to a prioritized text.

The alert criteria is described with reference to FIG. 14(a), which is a diagram of a user interface via which alert criteria options can be modified.

Referring now to FIG. 14(a), in text entry box 980, the user is able to enter the
10 email address of a pager or a cellular phone on which the user would like to be notified of prioritized text. Pagers and cellular phones are available that allow for paging by emailing text to an email address assigned to the pager or cellular phone. For example, as shown in FIG. 14(a), the device has an email address 4255555555@mobile.phoneco.net, where 4255555555 corresponds to the phone number of the cellular phone or the pager
15 number of the pager, as provided by the phone company "phoneco."

It is noted that the alert criteria of FIG. 14(a) specifically relates to new email messages. Three alert criteria are specifically shown in FIG. 14(a), and are referred to as options 982, 984 and 986. The options are not mutually exclusive, however. That is, the user can select one, two, or all three of the options 982, 984 and 986.

20 In the option 982, the user is able to specify that the electronic device should be contacted if a new email is received that has a priority greater than a predetermined threshold, and the user has been away from the computer for more than a predetermined amount of time. As shown in FIG. 14(a), the predetermined threshold is a priority of 85,

while the predetermined amount of time is 75 minutes. Thus, if it is determined that the priority of an email message is greater than 85, and that the user has been away from the computer for more than 75 minutes, then the user's electronic device will be alerted, consistent with the other options 984 and/or 986 if selected.

- 5 In the option 984, the user is able to specify that the electronic device should be contacted only when the current time is within a predetermined range of times. As shown in FIG. 14(a), the predetermined range of times is between 8:15 a.m. and 7:30 p.m. Thus, if it is determined that the current time is between 8:15 a.m. and 7:30 p.m., then the user's electronic device will be alerted, consistent with the other options 982 and/or 986 if
- 10 selected.

- In the option 986, the user is able to specific that if the user is in a meeting, then the user should only be notified if the priority is greater than a predetermined meeting threshold. In general, this meeting threshold is desirably greater than the threshold specified in the option 982. The purpose of this greater threshold is to indicate that the
- 15 user is receptive to notification of email messages greater than a particular priority in usual circumstances, but that the user is receptive to notification of email messages during meetings only if they have a priority that is unusually great. As shown in FIG. 14(a), then, the predetermined meeting threshold is a priority of 95, which is higher than the predetermined threshold of 85 in the option 982. Thus, if it is determined that the
- 20 user is in a meeting, and a message is received having a threshold greater than 95, then the user is alerted via his or her electronic device.

Routing Criteria

In this section of the detailed description, routing criteria are described. The routing criteria are the criteria that govern when and how a user is alerted to a prioritized text by having the text routed. The routing criteria is described with reference to FIGs.

- 5 14(b) and 14(c), which are diagrams of a user interface via which routing criteria options can be modified.

Referring first to FIG. 14(b), two tabs are selectable, a forward tab 1990, and a custom reply tab 1992. In FIG. 14(b), however, the forward tab 1990 is specifically selected, such that routing criteria with respect to forwarding a prioritized text such as an email message are shown. The user is able to specify an alternative email address to which high-priority emails are forwarded. More specifically, emails are forwarded to the address if the user has been away from the computer more than a predetermined amount of time, and a particular email to be forwarded has a priority greater than a predetermined threshold. For example, as shown in FIG. 14(b), the predetermined threshold is a priority of 95, and the predetermined amount of time is 600 minutes. Thus, if it is determined that the priority of an email message is greater than 95, and that the user has been away from the computer for more than 600 minutes, then the email will be forwarded to the specified email address.

Referring next to FIG. 14(c), the same two selectable tabs are shown, the forward tab 1990 and the custom reply tab 1992. However, in FIG. 14(c), the custom reply tab 1992 is specifically selected, such that routing criteria with respect to replying to the sender of a prioritized text such as an email message are shown. The user is able to specify a predetermined message that will be sent in the reply to the sender of the high-

priority email message. Emails are replied to if the user has been away from the computer more than a predetermined amount of time, and a particular email to be replied to has a priority greater than a predetermined threshold. For example, as shown in FIG. 14(c), the predetermined threshold is a priority of 95, and the predetermined amount of time is 120 minutes. Thus, if it is determined that the priority of an email message is greater than 95, and that the user has been away from the computer for more than 120 minutes, then the sender of the email will be replied to with the specified predetermined message.

10 Further Predetermined Criteria for Alerting

In this section of the detailed description, further predetermined criteria for alerting are described. The routing criteria is the criteria that governs when and how a user is alerted to a prioritized document, for example, by having the document displayed, or a sound played to indicate the arrival of the document. The criteria is described with reference to FIG. 14(d), which is a diagram of a user interface via which criteria options can be modified.

Referring to FIG. 14(d), the manner by which the user is alerted to a prioritized document depends on whether the document is classified as low, medium, or high priority. That is, depending on whether the document is classified as low, medium, or high priority, the alerting of the document is governed by the options selected in regions 2990, 2992 and 2994, respectively. The thresholds between a low and a medium priority document, and between a medium and a high priority document are user defined. Thus, a user indicates the threshold between low and medium in box 2996, and the threshold

between medium and high in box 2998. In other words, the user is alerting based on the priority being within a predetermined priority range. Alternatively, if the user selects check box 2991, the thresholds are set in a context-sensitive manner, by decision-theoretic reasoning.

- 5 Each region 2990, 2992, and 2994 has four options regarding the alerting of a prioritized document. First, the user may indicate that a sound is played when a document has a priority falling into a given region. Second, the user may indicate that a summary of the message and its priority receives focus, which means that a summary view is displayed listing the inferred priority of the message, and information
- 10 summarizing the content and nature of the message, including such information as the subject and sender of the message. This summary view may be temporarily selected as the active task on the screen, as opposed to the task currently being worked on by the user. Third, the user may indicate that the document be automatically opened in a window on the screen. Fourth, the user may indicate that an agent or automated assistant
- 15 with speech recognition and text-to-speech rendering abilities be activated, to alert the user to the priority and to allow the user to engage further in a dialog about hearing or seeing more about the message.

- Other alerting options include those selected by checking boxes 2993, 2995 or 2997. Checking box 2993 indicates that if the user is busy, then alerts are deferred until 20 the user is no longer busy, unless the priority of the document is at least a predetermined threshold. As shown in FIG. 14(d), this threshold is 75. Thus, it is first determined whether the user is busy and whether the priority of the document is greater than 75. If both of these conditions are true, then the user is alerted. Checking box 2995 indicates

that the window in which the document is being displayed is sized according to its priority, while checking box 2997 indicates that the window in which the document is being displayed is moved towards the center of the screen according to its priority.

5 Routing Priorities to a Mobile Device

FIGs. 15-20 are diagrams of a user interface for routing priorities to a mobile, or wireless, device different from that described in conjunction with FIG. 14(a). The mobile, or wireless, device can be a wireless phone, a personal digital assistant (PDA) having wireless communication capability, a pager, or another type of mobile device.

- 10 The user interface of FIGs. 15-20 presents the various options that a user can select from to have prioritized documents, which are also referred to as messages or items, routed to a mobile device.

- In FIG. 15, the user interface 4500 includes a summary window 4502. The window 4502 indicates the currently active profile 4504, the wireless service provider 15 4506, and the device statistics 4508. As shown in FIG. 15, the profile 4504 is "out of office," indicating the prioritized documents are routed to the wireless device based on the "out of office" profile. The provider 4506 includes a logo 4510, which can be clicked on to transport the user to the web site of the provider 4506. The statistics 4508 indicate the active wireless device 4512, and the number of documents, or items, 4514 that have 20 been sent since the last time the statistics were reset. The statistics can be reset by pressing the button 4516.

The left pane 4518 of the user interface 4500 is divided into three areas, preferences 4520, a profile editor 4522, and device configuration 4524. As shown in

FIG. 16, selecting the “My Schedule” item 4526 of the preferences 4520 brings up the window 4528. The window 4528 indicates when priorities should be transmitted to the wireless device based on the work profile 4530, and when the priorities should be transmitted based on the home profile 4532. As indicated in FIG. 16, the user selects the various days and times within area 4534 as to when the work profile 4530 should be used, such that the home profile 4532 is used for the days and times not selected.

- 5 As shown in FIG. 17, selecting the “My Daily Summary” item 4536 of the preferences 4520 brings up the window 4538. The window 4538 indicates when calendar summaries should be sent to the wireless device. A calendar summary is a summary of the user’s appointments in his or her calendar for a given day. The user specifies which days of the week in the area 4540 and at what time of day in the area 4542 the summaries should be sent. The summaries sent can be for the current day, or the next working day, as chosen by the user in the area 4544. The user can also force transmission of the calendar summary to the wireless device by pressing the button 4546.
- 10

- 15 As shown in FIG. 18, selecting the work profile 4530 from the profile editor 4522 brings up the windows 4548, 4550, 4552, and 4554. The work profile 4530 is one of the profiles that can be edited within the profile editor 4522. Having multiple profiles enables the user to have different schemes for priority transmission based on different contexts. For example, the user may have one type of profile when he or she is at work, another when at home, another when on vacation, and so on. Profiles can be as granular as desired by the user. For example, the user may have a different profile for each day of the week, or one profile for weekdays and another for weekends.
- 20

Clicking the button 4556 in the window 4548 enables the user to change the schedule for the work profile 4530. The user can also lock the profile by checking the check box 4558. Locking the profile causes it to remain as the active profile until it is unlocked. The window 4550 allows the user to specify how email should be prioritized

5 and sent to the wireless device. Selecting one of the options in the area 4560 allows the user to indicate whether more or less email, as sorted by priority, should be sent to the device. Furthermore, the user can select the button 4562 to specify rules as to how messages can also be specified for transmission to the device. That is, the options in the area 4560 indicate the number of emails, as sorted by priority, that should be sent to the

10 device, whereas selecting the button 4562 allows specification of rules as to how emails should also be sent to the device.

The different options that the user can select from in the area 4560 is a coarsening mechanism designed to make the editing of a profile easier for the user. For example, as shown in the area 4560, there are five different options that the user can select, from

15 sending no messages, to sending all messages. The options in between specify that only high priority messages should be sent, only high and medium priority messages should be sent, and only low, high and medium priority messages should be sent. The coarsening mechanism is compared to allowing the user to set a precise priority score threshold, such that messages having priorities over the threshold are sent.

20 The window 4552 enables the user to specify whether reminders regarding events in the user's calendar, or tasks in the user's task list, should be sent to the wireless device. By checking the check box 4564, the user can specify that such reminders be sent only after the user's computer has been idle for a user-specified length of time. The length of

time set in the window 4552 is an idle threshold. Finally, the window 4554 allows the user to specify when priorities are related to the wireless device. For example, in the area 4566, the user can specify that messages be delivered only every x number of minutes, where x is user specified, or until there are y number of messages to send, where y is also user specified. In the area 4568, the user can specify that messages be delivered only after the computer has been idle for a user-specified length of time.

The ability of the user to specify when priorities are delivered in the areas 4566 and 4568 is referred to as priority chunking. Priorities are “chunked” into groups by time, by number of messages, or both, before the user is alerted to them on the wireless device. Chunking may also be accomplished by priority density. That is, a user receives alerts on the mobile device when the total priority of the chunked messages exceeds a threshold. Utility-based transmission and caching can also be accomplished. In general, both compression and transmission can be selective as to certain messages.

As shown in FIG. 19, selecting the message formatting item 4570 of the device configuration 4524 brings up the window 4572. The window 4572 allows the user to control how messages are formatted when they are sent to the wireless device. The user can specify in the area 4574 whether messages should be compressed, and if so, how compressed they should be. The area 4574 provides a coarsening mechanism for specifying compression, similar to that described in conjunction with the area 4550 of FIG. 18. The user can also specify in the area 4576 whether fragmentation should be enabled, according to one of two settings. A sample message in the window 4578 is shown as formatted in the window 4580, according to the settings specified in the areas 4574 and 4576. Formatting, such as compression, fragmentation, and truncation, may

also be automated. For example, higher-priority messages may be compressed less than lower-priority messages.

- Finally, as shown in FIG. 20, selecting the device selection item 4582 of the device configuration 4524 brings up the window 4584. The window 4584 allows the user 5 to select the type of wireless device to which priorities are sent. In particular, the user specifies the type of wireless device by manufacturer, as indicated in the area 4586, and model, as indicated in the area 4588.

Extensions

- 10 Extensions to the routing of priorities, as described in the previous sections, can also be accomplished. One extension is to make available tags within the user interface that capture ontology of appointment and meeting types by degree of interruptability. This allows users to mark events within the calendar with these tags, or have them automatically marked based on the content of the meeting. The tags can then be used to 15 control policies for sending priorities to the mobile device. For example, there can be the following tags: no meeting, normal meeting, low-threshold interruptable meeting, and high-threshold interruptable meeting. When there is no meeting, or a normal meeting, the mobile device can normally receive priorities. However, when there is a low-threshold or a high-threshold interruptable meeting, the mobile device only receives priorities meeting 20 the low or high threshold.

Another extension is the detection of conversation at the desktop or at the wireless device. If conversation is detected, then the user will not be alerted to priorities until the

conversation ceases. Furthermore, the user may specify not to receive alerts between certain times, unless the alert is of sufficiently high priority.

Conclusion

5 Although specific embodiments have been illustrated and described herein, it will
be appreciated by those of ordinary skill in the art that any arrangement that is calculated
to achieve the same purpose may be substituted for the specific embodiments shown.

This application is intended to cover any adaptations or variations of the present
invention. Therefore, it is manifestly intended that this invention be limited only by the
10 following claims and equivalents thereof.

I claim:

1. A method operable on a text comprising:

generating a priority of the text based on a classifier by determining a likelihood

that the text is of high priority;

- 5 determining an expected loss of non-review of the text at a current time based on the priority;

determining an expected cost of outputting the text at the current time; and,

alerting a user of the text on a mobile device of the user in response to

determining that the expected loss is greater than the expected cost, in accordance with a

- 10 current profile.

2. The method of claim 1, wherein determining the expected loss of non-review

comprises determining a likelihood that the user will review the text at a future time.

3. The method of claim 1, wherein determining the expected loss of non-review

comprises determining a current expected rate of lost opportunity for the user resulting

- 15 from non-review of the text as a function of time.

4. The method of claim 1, wherein the text classifier is one of a Bayesian classifier

and a support-vector machine classifier.

5. The method of claim 1, wherein the current profile is one of a plurality of profiles,

each of the plurality of profiles editable by the user to reflect a different context.

6. The method of claim 5, wherein each of the plurality of profiles is schedulable on a per-day and by-time basis.
7. The method of claim 5, wherein each of the plurality of profiles is lockable to remain active unless unlocked.
- 5 8. The method of claim 5, wherein each of the plurality of profiles has a priority threshold settable by the user through a coarsening mechanism.
9. The method of claim 5, wherein each of the plurality of profiles has a chunk setting such that the text is delivered to the mobile device only in conjunction with one or more other texts.
- 10 10. The method of claim 5, wherein each of the plurality of profiles has a chunk setting such that the text is delivered to the mobile device only when a specified period has expired.
11. The method of claim 5, wherein each of the plurality of profiles has a setting such that the text is delivered to the mobile device only if a primary device of the user is idle
- 15 for a specified period of time.
12. The method of claim 1, further comprising, prior to alerting the user on the mobile device, formatting the text.

13. The method of claim 12, wherein formatting the text comprises compressing the text according to a specified compression setting.
14. The method of claim 13, wherein the specified compression setting is settable by the user through a coarsening mechanism.
- 5 15. The method of claim 12, wherein formatting the text comprises fragmenting the text according to a specified fragmentation setting.
16. The method of claim 1, further comprising sending to a sender of the text indication that the text has been delivered to the mobile device of the user.
- 10 17. The method of claim 1, further comprising sending to a sender of the text indication that the user is away from a primary device of the user.
18. The method of claim 1, wherein the classifier considers a structural relationship between the user and a sender of the text.
19. The method of claim 1, wherein the classifier considers a time of an event referenced in the text.
- 15 20. The method of claim 1, wherein alerting the user on the mobile device is further in accordance with a calendar of the user, the calendar having one or more events, each event associated with a tag specifying an interruptability degree of the event.

21. The method of claim 1, wherein the method is performed by execution of a computer program by a processor from a computer-readable medium.
22. A method operable on a text comprising:
- generating a priority of the text based on a classifier; and,
- alerting a user of the text on a mobile device of the user in accordance with a current profile, the current profile selected from a plurality of editable profiles, each of the plurality of profiles schedulable on a per-day and by-time basis and having a priority threshold settable by the user through a coarsening mechanism and a chunk setting controlling delivery of the text.
- 10 23. The method of claim 22, wherein the chunk setting is such that the text is delivered to the mobile device only in conjunction with one or more other texts.
24. The method of claim 22, wherein the chunk setting is such that the text is delivered to the mobile device only when a specified period has expired.
- 15 25. The method of claim 22, wherein each of the plurality of profiles further has a setting such that the text is delivered to the mobile device only if a primary device of the user is idle for a specified period of time.
26. The method of claim 22, further comprising sending to a sender of the text indication that the text has been delivered to the mobile device of the user.

27. The method of claim 22, further comprising sending to a sender of the text indication that the user is away from a primary device of the user.
28. The method of claim 22, wherein alerting the user on the mobile device is further in accordance with a calendar of the user, the calendar having one or more events, each 5 event associated with a tag specifying an interruptability degree of the event.
29. The method of claim 22, wherein the method is performed by execution of a computer program by a processor from a computer-readable medium.
30. A method operable on a text comprising:
generating a priority of the text based on a classifier;
10 formatting the text according to a specified compression setting settable by the user through a coarsening mechanism and according to a specified fragmentation setting;
alerting a user of the text as formatted on a mobile device of the user in accordance with a current profile.
31. The method of claim 30, wherein the current profile is selected from a plurality of 15 editable profiles, each of the plurality of profiles schedulable on a per-day and by-time basis and having a priority threshold settable by the user through a coarsening mechanism and a chunk setting controlling delivery of the text.
32. The method of claim 30, further comprising sending to a sender of the text indication that the text has been delivered to the mobile device of the user.

33. The method of claim 30, further comprising sending to a sender of the text indication that the user is away from a primary device of the user.
34. The method of claim 30, wherein alerting the user on the mobile device is further in accordance with a calendar of the user, the calendar having one or more events, each event associated with a tag specifying an interruptability degree of the event.
- 5
35. The method of claim 30, wherein the method is performed by execution of a computer program by a processor from a computer-readable medium.


SEARCHED - INDEXED - SERIALIZED - FILED

ABSTRACT OF THE DISCLOSURE

- Generation and management of priorities are disclosed. A priority for a text is generated using a classifier. The text may be formatted, such as compressed or fragmented. The text is delivered to a mobile device based on a profile. The profile is a user-editable and schedulable profile. The profile has a priority threshold settable by the user through a coarsening mechanism, and a chunk setting controlling delivery of the text. The text may be delivered to the mobile device only if a primary device of the user is idle for a period of time. The sender of the text may be sent indication that the user is away from the primary device, or that the text has been delivered to the mobile device.
- Alerting the user may consider the user's calendar, each calendar event having a tag specifying the degree to which the user is amenable to being interrupted.

I hereby certify that this is being deposited with the United States Postal Service "Express Mail Post Office to addressee" service under 37 CFR § 1.10 in an envelope addressed to The Assistant Commissioner for Patents, Washington, DC 20231, on Dec 8, 2000, by Michael Dryja, and having "express mail" mailing label no.

EL045149846US.



Signature of Michael Dryja

THIS PAGE BLANK (USPTO)